

# What makes an Image Iconic? A Fine-Grained Case Study

Yangmuzi Zhang<sup>†</sup>, Diane Larlus<sup>\*</sup>, and Florent Perronnin<sup>\*</sup>

<sup>†</sup>University of Maryland, College Park, USA

<sup>\*</sup>Xerox Research Centre Europe (XRCE), Meylan, France

**Abstract.** A natural approach to teaching a visual concept, *e.g.* a bird species, is to show relevant images. However, not all relevant images represent a concept equally well. In other words, they are not necessarily iconic. This observation raises three questions. Is iconicity a subjective property? If not, can we predict iconicity? And what exactly makes an image iconic? We provide answers to these questions through an extensive experimental study on a challenging fine-grained dataset of birds. We first show that iconicity ratings are consistent across individuals, even when they are not domain experts, thus demonstrating that iconicity is not purely subjective. We then consider an exhaustive list of properties that are intuitively related to iconicity and measure their correlation with these iconicity ratings. We combine them to predict iconicity of new unseen images. We also propose a direct iconicity predictor that is discriminatively trained with iconicity ratings. By combining both systems, we get an iconicity prediction that approaches human performance.

## 1 Introduction

Humans often associate a concept, *e.g.* an object, a scene, a place or a sentiment, with a normalized visual representation, referred to as a canonical representation. This observation motivated the introduction of the notion of a *canonical* or *iconic* image: an image is said to be canonical/iconic w.r.t a given concept if it is a “good representative” for the said concept. Rosch and Palmer [29] showed in their seminal work that humans agree on canonical views of objects, and that recognition is faster for these views.

Several works have considered the task of predicting image *iconicity* [2,3,11] [19,24,26,30,37]. Such iconicity predictors have many applications. In the graphics domain, this can be used to choose the best illustration of a concept [15]. For image search on the web, iconicity can be exploited to rerank the top-retrieved results [19,30]. In the case of a consumer photo application, it allows summarizing a collection of photographs, *e.g.* a set of holiday pictures [33,37]. Finally, for a semi-automatic visual annotation system which involves a human in the recognition loop [5,35], iconicity prediction enables choosing the “best” images to display, *i.e.* those that will help annotators making their decision. We are particularly interested in this last scenario, and more precisely in difficult visual



**Fig. 1.** Given the set of images on the left, which one would you use to teach what a *Barn Swallow* is? The picture with the green check mark is selected automatically by our method. On the contrary, the picture with the red cross is regarded as the least suitable. The rationale for these decisions is explained in Fig. 3.

recognition tasks that require expert knowledge but are crowd-sourced to non-expert annotators. This includes fine-grained recognition tasks which involve a large number of visually similar and semantically related classes (*e.g.* species of birds and flowers, brands and makes of vehicles, etc.). In such a case, annotators cannot solely rely on class names (*e.g.* 'Barn Swallow') or descriptions that are very technical. They need to be provided with appropriate – iconic – visual representations. To our knowledge, the question of how to choose such iconic images for annotation interfaces has been largely overlooked by the computer vision community. In this work, we raise three questions. *Is iconicity a purely subjective property? Can we predict an image's iconicity with respect to a concept? What makes an image iconic?* We provide answers in a fine-grained classification scenario.

The first requirement for our study is a dataset. Since our focus is on fine-grained tasks and we are not aware of any publicly-available fine-grained dataset with iconicity scores, we leverage the CUB dataset [34] which contains images of 200 bird species. We collected iconicity scores from non-expert annotators. In line with our primary target application – recognition with a human in the loop – they were asked to choose the images that they would show to teach what a bird species looks like. We note that the level of expertise of the annotator is likely to play a role in rating image iconicity. While studying what makes an image iconic for experts would also be interesting, we believe there is much value in studying this problem when the users are non-experts since the majority of annotators encountered on crowdsourcing platforms are not.

The second requirement is to establish a large list of properties that may play a role in deciding whether an image is iconic. This includes the object size and position, the visibility of its parts or the presence of the attributes associated with the class, the aesthetics and the memorability of the image, the similarity of this image to the average representation of its class, and its discriminability w.r.t. its class. These properties are quantified by what we later refer to as *iconicity indicators*. We also describe approaches to measuring these indicators. We measure their correlation with iconicity. This is in contrast with the vast majority of the previous studies on the topic which usually focus on a single property, for instance the viewpoint (see section 2 for a review of related work), or provide only qualitative results.

Finally, we consider iconicity predictors directly trained on generic image descriptors using iconicity labels. To the best of our knowledge, such a direct approach to iconicity prediction has never been considered. Despite its simplicity, we show that it yields surprisingly competitive results and that it is complementary to the indicator-based approach.

In summary, the contributions of this work are manifold. First, we enrich an existing public dataset with iconicity ratings and show that agreement between annotator ratings is significant, *i.e.* that iconicity is not a purely subjective property. Second, we propose an extensive list of properties likely to be relevant to predict iconicity and measure the correlation between these properties and the user iconicity ratings. To the best of our knowledge, this is the first quantitative study of its kind. Third, we study the combination of these different properties to predict the iconicity of a new image. Fourth, we propose a direct approach to iconicity prediction and show that it obtains competitive results. Combining all our predictors we achieve a prediction accuracy that approaches the human performance upper-bound. Finally, we provide qualitative results showing that the conclusions we draw from bird images can be extrapolated to very different types of objects, such as planes and shoes.

## 2 Related Work

In this section, we review the properties that have been used in previous works to predict iconicity. Mainly two properties have been considered: the viewpoint or the ability to summarize a collection. Only few works considered properties beyond viewpoint and summarization, and combined these properties. We also outline the limitations of previous evaluations which were mostly qualitative.

**Iconicity and viewpoint.** Following the seminal work of [29] that showed evidence of a consistently preferred viewpoint, many works studied the link between iconicity and viewpoint, and considered different photos of the same object instance, typically viewed under ideal conditions (*e.g.* synthesized object with no background). Several user studies have verified the existence of iconic viewpoints for 3D objects [4,6] as well as for scenes [10]. Several works have also considered the problem of computing the best viewpoint from a 3D model [36] or a set of 2D shapes [9]. We agree that the viewpoint plays an important role in finding a good representative for a category, but argue that other properties should also be taken into account in the definition of iconicity in more realistic scenarios.

**Iconicity and summarization.** Many works also considered the case where the image set is a large collection of noisy images collected from the Internet, for instance by querying a search engine such as Google Image Search or a photosharing website such as Flickr for a specific concept. In this scenario, an iconic image is an image that best summarizes the data and the problem of finding iconic images is generally treated as one of finding clusters [19,24,30] or modes [26,37] in the image feature space. In most of these works, the results are evaluated either qualitatively through a manual inspection of the found iconic images [26,30,37] or simply by measuring whether the found iconic images are

relevant or not with respect to the concept [24]. Yet, according to our definition, a relevant image is not necessarily iconic. Only Jing et al. [19] conducted a user study to evaluate whether the found iconic images were preferred to random images.

**Beyond viewpoint and summarization.** Berg and Forsyth [3] proposed a nearest-neighbor classifier to predict image iconicity and used figure/ground segmentation to focus on the foreground. Images of landmarks collected on Flickr were evaluated by 23 volunteers as representing well the landmark or not. However, their study focuses on instances, not categories, and does not provide any detailed analysis as to what makes an image iconic. Berg and Berg [2] suggested properties that could correlate with iconicity such as the object size and position. However, in their evaluation, the users were explicitly instructed to take these criteria into account which biased the results somewhat favorably toward these properties. Raguram and Lazebnik [30] proposed to leverage an aesthetic measure but only a qualitative evaluation of the impact of the aesthetic factor was conducted. Ehinger et al [11] showed that images typical for scene categories were classified correctly with high confidence. Our study considers properties suggested in [2,3,30] but also new ones, and reports a detailed analysis of the correlation between each property and iconicity. We also combine complementary indicators and propose a full quantitative study. Note that direct iconicity prediction has never been considered in the past.

**Beyond naming.** While much work in the computer vision community has been devoted to *naming* objects and scenes, more and more recent works have proposed to *describe* them according to their parts and attributes [14,22,12], their aesthetic value [21,27] or their memorability [18]. Our quantitative evaluation of iconicity fits in this line of research as it goes beyond naming objects and scenes. Also, we measure how these properties – attributes, aesthetics, memorability – correlate with iconicity and use them as indicators in our study.

### 3 Dataset

We base our study on the Caltech-UCSD-Birds-200-2011 dataset (CUB for short) [34], which contains 11,788 images of 200 bird species<sup>1</sup>. We chose CUB for the following reasons. First, bird species recognition is a fine-grained task, and is extremely challenging both for computers and non-expert annotators. For such problems, hybrid annotation systems that involve a user in the recognition loop [5,35] have been explored. As explained earlier, non-expert annotators cannot solely rely on the bird names or descriptions to make decisions. Choosing the proper iconic images is likely to make a vivid difference in such hybrid systems. Second, the dataset contains realistic images of birds in the wild. Although images usually have good resolution, by far not all of them can be considered iconic: see Fig. 1, 2 and 3. Third, this dataset comes with a rich set of annotations: bounding boxes locate the birds precisely in images, all parts relevant to

---

<sup>1</sup> In this paper we use the words “species” and “classes” interchangeably.



**Fig. 2.** Annotators rated the iconicity of images in sets of 5.

birds (*e.g.* beak, eyes, legs, etc.) are indicated as visible or not, and image-level attribute annotations describe which visual attributes can be observed in each image. This allows an in-depth study of a large set of iconicity indicators. We collected iconicity ratings to enrich this dataset<sup>2</sup> based on the following protocol.

**Acquiring iconicity scores.** We acquired annotations from a set of 32 non-expert annotators. Each participant was shown 50 sets of images corresponding to 50 bird classes and asked to rate the iconicity of each image, where an iconic image was defined as *the kind of images one would use to show a person what a particular bird species looks like*, in agreement with the definition we give in the introduction. The iconicity could be rated according to three values: 0 for “bad”, 1 for “fair”, and 2 for “good”. For each class, 5 images of the same bird species were shown (in one row, with the same height, see Fig. 2). Hence, even if the participants were not acquainted with a particular bird species, showing multiple images simultaneously provided them with an opportunity to get familiarized.

The dataset was initially divided by [34] into 5,994 training images, and 5,794 test images. For the training set, we collected annotations from 20 participants, following a split of the data that uses some redundancy (see next paragraph), so we obtained 4,100 iconicity annotations for the training set. The annotated test set is composed of 2,995 images which were annotated by a set of 12 users<sup>3</sup>. *To avoid any bias, we made sure that participants annotating the training and test sets were strictly different.*

We underline that *iconicity is not an absolute property but a relative one*: the most iconic image of a set depends on the other images in the set. Especially, an image might be judged more iconic with respect to a bird class if it is compared to a set of random bird shots and less iconic if it is compared to images which have been uploaded on Flickr (as is the case of the CUB images). To take this fact into account, our evaluation considers relative values, *i.e.* ranks. This disregards global shift or scaling effects on the scores.

**Consistency among annotators.** Iconicity could be seen as a subjective property. Hence, we first conduct an experiment to measure inter-person agreement. For that purpose, during the acquisition of the training set, we divided the 20 participants into 2 groups. All the persons in one group had 50 images (from 10 different classes) in common. We calculate the correlation between annotations for every pair of annotators from the same group.

<sup>2</sup> iconicity annotations available upon request

<sup>3</sup> 15 images were annotated for each of the 200 species, except for 2 classes that have only respectively 11 and 14 images in the test set.

The correlation is measured using the Spearman’s Rank Correlation (SRC) coefficient, which can handle rank ties. The  $n$  scores  $A_i$  and  $B_i$  are converted to ranks  $a_i$ ,  $b_i$ , and SRC is computed as:

$$\rho = \frac{\sum_{i=1}^n (a_i - \bar{a})(b_i - \bar{b})}{\sqrt{\sum_{i=1}^n (a_i - \bar{a})^2 \sum_{i=1}^n (b_i - \bar{b})^2}} \in [-1, 1] \quad (1)$$

where  $\bar{a}$  and  $\bar{b}$  are the mean of the ranks  $a_i$  and  $b_i$  respectively. The sign of  $\rho$  indicates the direction and strength of association between the score sets  $A$  and  $B$ , where 1 means that the two ranks perfectly match, 0 no correlation, and -1 that they are anti-correlated. The corresponding p-values are also calculated. If the p-value is small, *e.g.* less than 0.05, it means the correlation between annotations is significantly different from 0 with a 5% confidence level [7]. For groups 1 and 2, we measured SRCs of 0.485 and 0.497 with p-values of 0.045 and 0.006 respectively. Both p values are lower than 0.05. Hence, despite the subjective nature of iconicity, *there is a strong agreement between annotators*. This shows that *iconicity is not purely subjective, even when annotators are not domain experts*. Therefore, we limited ourselves to a single iconicity score recorded per image, in order to cover as much of the CUB dataset as possible.

**A few statistics.** Among the training images, 1,597 images (39%) are rated 2, 1,742 images (43%) are rated 1, and 761 images (19%) are rated 0. The testing set follows the same trend, with 1,161 images (39%) rated 2, 1,257 images (42%) rated 1, and the rest rated 0. This shows that the full iconicity scale was used.

## 4 Measuring Iconicity

We select several properties that we expect to correlate with image iconicity and quantify them using a variety of indicators. Some of these properties have been used in previous works. We also propose new ones such as the ones based on attributes, occlusion or memorability. Each indicator produces a score for each image. Some of these indicators based on the available ground-truth annotations (*e.g.* a bounding box) are referred to as oracles. We leverage the rich set of annotations available with CUB for this purpose (see previous section). When relevant, we also consider alternative indicators that use predicted information (*e.g.* from an object detector) instead of the ground-truth. These predicted indicators can be compared to their oracle counterparts. The properties are divided into i) class-independent indicators that do not need to know which class an image belongs to and ii) class-dependent indicators that use the class label. We finally consider the task of predicting iconicity.

### 4.1 Class-independent indicators

**Object size and location.** We first look at simple statistics capturing the scene composition. As in [2], we look at the object size (iconic images are expected to present a large object) and location in the image, using the ground-truth

bounding-box (BB) bird location. We derive two indicators: **BB-size** measures the percentage of image pixels covered by BB, and **BB-dist2center** computes the distance between the object center (BB center) and the image center, normalized by the length of the image diagonal. We also study the case where the object location is unknown, and use the state-of-the-art Deformable Part Model (DPM) object detector [13,16]. The DPM is trained using the BB annotations of the 200 species of birds from the training set to build a generic bird detector. We define two indicators computed using the DPM output instead of the ground-truth, **DPM-size** and **DPM-dist2center**.

**Occlusion.** Images in CUB are annotated with the location of 15 bird body-parts<sup>4</sup>. For each part, we know whether it is visible or not. We use this information as an occlusion indicator, where the **Occlusion** score is simply the number of visible parts. Although view-point seems to have played a crucial role in several previous studies on iconicity in constrained scenarios (see related work), it is unclear how to define such a criterion when dealing with realistic images, and especially with articulated objects, as is our case. Indeed if the body of the bird faces one direction, the head can face another. Instead of building a view-point estimator that would be ill-defined in our case and probably of low accuracy, we use the occlusion criterion as a proxy.

**Aesthetic scores.** Although aesthetics and iconicity are not explicitly related, we expect images of high aesthetic quality to have a higher chance to be considered iconic. This is because visually pleasing images are generally of high quality and well-composed [21], and aesthetic criteria influence choosing a representative for teaching purposes (our scenario). This intuition had already been proposed by [30], but only evaluated qualitatively. We evaluate it quantitatively in our work. Our aesthetic predictor is based on [27], which trains a classifier directly on a patch-based image representation. This approach was shown to implicitly capture the photographic rules which are explicitly encoded by [21], while providing a superior performance. As suggested in [27], we use Fisher Vector (FV) image representations [32] (see section 5 for more details). To train our aesthetic indicators, we leverage the large AVA dataset [27] which contains  $\sim 200,000$  images labeled with binary aesthetic labels (high or low aesthetic quality). We first consider a generic model, trained with the full training set. This model is then applied to the bird images, and the score of an image is simply the SVM score. This corresponds to indicator **Aesthetic-Generic**. A subset of images of the AVA dataset is also annotated with semantic tags. The most relevant tag to birds is “animal” ( $\sim 2,500$  images). Therefore, we also trained an animal-specific aesthetic model. This indicator is referred to as **Aesthetic-Animal**.

**Memorability.** Memorability measures how well an image can be remembered by a person. We hypothesize that memorability and iconicity share common properties. [17,18] showed that image memorability prediction is possible with current computer vision techniques. Consequently, we train a memorability predictor using the SUN Memorability dataset [17], that contains 2,222 images

---

<sup>4</sup> Parts are: both eyes, the forehead, the crown, the bill, the nape, the throat, the breast, the back, both wings, the belly, both legs and the tail.

labeled with memorability scores. Again, we use Fisher Vector (FV) representations [32]<sup>5</sup>. The memorability scores from the SUN dataset are used to train a linear SVM classifier. The SVM scores on the test image constitute our **Memorability** predictor. The dataset contains only few images with animals (less than 10) so we did not train an animal-specific memorability model.

## 4.2 Class-dependent indicators

We now consider as iconicity indicators for a given image those properties which quantify the relevance of the image with respect to its class. We therefore assume a labeled training set of  $n$  images  $L = \{(x_1, y_1), (x_2, y_2), \dots, (x_n, y_n)\}$ , where  $x_i$  is the feature vector of image  $i$  and  $y_i$  is its label, where  $y_i \in \{1, \dots, K\}$  and  $K$  is the number of classes ( $K = 200$  for CUB).

**Distances to the cluster center.** As mentioned in section 2, most previous works treated the problem of finding iconic images as one of finding clusters or modes. Our first set of indicators follows related work. Since we have very few training images per class (on the order of 30), it is unreasonable to run a clustering algorithm per class. Therefore, we assume that each class  $k$  has a single cluster whose mean is denoted  $\mu_k$ . For any new image  $x'$  from the test set, we compute a similarity to the cluster center ( $-||x' - \mu_k||^2$ ), and use this score as iconicity measure for class  $k$ . Following [24,31], we first consider the GIST descriptor [28] as an image feature vector. This indicator is denoted **Cluster-GIST**. The average GIST descriptor of a class was also used as a semantic representative in [30,24]. We also consider the Fisher-Vector (FV) [32] representation which results in the **Cluster-FV** indicator.

**Object classifier scores.** An image that receives a high score from a classifier trained to recognize one class should represent this class well because it is supposed to contain more discriminative features. Consequently, we consider using classifier scores as one of our indicators. This was also used to measure scene typicality of images in [11]. More precisely, we train one linear SVM classifier per class using the labeled training set  $L$ . Then the iconicity of a new image  $x'$  with respect to a class can be measured by computing the corresponding classifier score. Again, we consider two types of descriptors: the GIST descriptor [28] (indicator denoted **SVM-GIST**) and the Fisher Vector [32] (**SVM-FV**).

**Classifiers using attributes.** CUB contains annotations for  $M=312$  attributes<sup>6</sup>, at the image level. In other words, each image is associated with an attribute representation  $a = [a_1, a_2, \dots, a_M]$ , where  $a_i$  takes binary values to indicate whether an attribute is present in this image or not. All the remaining indicators are based on the intuition that an iconic image for a given class should best display the attributes of that class. To the best of our knowledge, this is the first time that this criterion has been considered and evaluated for image iconicity. We considered 4 different indicators based on attributes.

<sup>5</sup> To validate our feature choice, we reran the experiment in [18] and obtained a correlation similar to the one they report (without the semantic attributes).

<sup>6</sup> The 312 attributes describe the bird color (of the wings, the back, the forehead, etc.) and shape (of the bill, the tail, etc.).



First, we can use these image-level attribute annotations together with a distance-based classifier. Let us assume that we have a class-level attribute vector:  $c = [c_1, c_2, \dots, c_N]$  (built by averaging image-level attribute vectors). We define an image-to-class similarity (I2C) between an image and class  $C$  as:  $-||a - c||$ . This similarity is used as our indicator score and is referred to as **I2C-Att-Orac**.

We also use the image-level attribute vectors as image representations, and we directly train SVM classifiers on top to recognize bird species, using the attribute vectors of the training set. Then the trained per-class classifier can be used to predict a score for each image. We denote this indicator **SVM-Att-Orac**.

Finally, the DAP [22] model is a standard way to predict categories based on attribute-level information. Given image  $x$ , we first obtain the attribute predictions  $p(a_m|x)$  by training  $M$  independent attribute classifiers. The score of image  $x$  is then given by  $p(a|x) = \prod_{m=1}^M p(a_m|x)$ . We considered an oracle scenario and a prediction one. For the oracle scenario, the probabilities are 0 or 1, based on the image level annotation. We use  $\epsilon=10^{-5}$  for probability 0 and  $1-\epsilon$  for probability 1 to avoid the overall  $p(a|x) = 0$ . This indicator is referred to as **DAP-Orac**. For the prediction scenario, we assume that test images are not annotated with image-level attribute vectors, and we predict attribute probabilities using attribute classifiers trained on training images (**DAP-Pred**). All DAP models are learnt on FV representations.

### 4.3 Iconicity Prediction

Given the previous indicators, we can now predict iconicity. The simplest approach consists in linearly combining the indicators, either by giving them equal weights, or by learning a vector of weights  $w$  using iconicity labels. In the latter case, we concatenate the indicator values, whiten the resulting feature vector and learn a linear SVM. We experimented with two learning frameworks: a binary SVM classifier and a ranking SVM [20]. A disadvantage of the binary SVM formulation is that it requires setting an arbitrary threshold that will split the training set into iconic *vs.* non-iconic images. In our experiments, we set this threshold to 1.5. The ranking SVM formulation deals directly with ranked training pairs, which fits better the relative nature of iconicity. More formally, given training pairs of images  $(x^+, x^-)$  such that  $x^+$  is ranked higher than  $x^-$ , we minimize the regularized ranking loss:

$$\sum_{(x^+, x^-)} \max\{0, 1 - w'(x^+ - x^-)\} + \frac{\lambda}{2} ||w||^2. \quad (2)$$

Note that, in our implementation, the images  $x^+$  and  $x^-$  of a given pair  $(x^+, x^-)$  are from the same batch of 5 images annotated by the same user. This is to avoid confounding factors during the learning process (*e.g.* the fact that a user might be more inclined to rate images as iconic than other users).

As an alternative to the indicator-based approach to iconicity prediction, we also consider the approach which consists *in predicting the iconicity directly from image representations such as the FV*. We refer to this approach as **Direct**

Iconicity Prediction (DIP). As is the case for the indicator-based predictor, we use a linear SVM on the FV features and its parameters can be learned using either a binary classification objective function or a ranking objective function. To our knowledge, this is the first attempt to such a direct approach to iconicity prediction. These predictors are denoted **DIP-bin** and **DIP-rank**

## 5 Experiments

After discussing implementation details (section 5.1), we first evaluate the correlation of each indicator with the iconicity ratings and with each other (section 5.2). In the second set of experiments, we evaluate how well we predict iconicity by combining different indicators and by direct prediction (section 5.3).

In all experiments we use the standard train/test split of CUB. All supervised learning (*e.g.* training DPMs, computing class means, training SVM classifiers) is performed on the train split and all the results are measured on the portion of the test split which was rated by the users. For validation, we split the train set into two halves: we train on the first half, validate on the second half and retrain the models on the full training set with the optimal validation parameters.

### 5.1 Implementation details and evaluation

**Implementation details.** We use GIST and Fisher-Vector (FV) descriptors in some of the listed indicators. For the GIST, we used the color implementation of [28] (960 dimensions). We compute FV representations on top of SIFT descriptors [25] and color descriptors [8] for each image. The number of Gaussians is 1,024. We use a spatial pyramid [23] with 8 regions: the full image, 3 horizontal stripes and the 4 quadrants. There is one FV pipeline for each low-level descriptor, and both pipelines are combined with late fusion (score averaging).

**Measures.** We consider two different evaluation measures. First, in accordance with the binary classification setting, we evaluate the quality of indicators using the average precision (AP). We define as positive (= iconic) those images whose ground-truth label is above a threshold  $\alpha=1.5$  and the remaining are labeled as negative. Second, in accordance with the ranking setting, we look at the correlation between indicators and iconicity. Again, it is measured using Spearman’s Rank Correlation (SRC) coefficient (see eq. (1)). SRC is computed between the indicator scores and the ground-truth iconicity scores. The corresponding  $p$ -values are also reported. A  $p$ -value smaller than 0.05 means statistical correlation with a 95% confidence level.

### 5.2 Analysis of iconic images

We first evaluate the correlation of each indicator with the iconicity ratings and with each other. Results are presented in Table 1. At first glance, all the methods exhibit a correlation with iconicity, except **SVM-GIST** ( $p$ -value  $> 0.05$ ). Although we can observe small differences between the ranking of the methods between

**Table 1.** Comparison of the proposed indicators on the test images. Spearman rank correlation (SRC), corresponding p-value, and average precision (AP) are reported.

Class-independent Indicators				Class-dependent Indicators			
Strategy	SRC	p-value	AP	Strategy	SRC	p-value	AP
BB-size	0.304	3.01e-65	51.7	Cluster-GIST	0.048	9.20e-03	41.7
DPM-size	0.280	6.84e-55	51.6	Cluster-FV	0.111	1.18e-09	43.3
BB-dist2center	0.097	9.53e-08	43.4	SVM-GIST	0.027	1.44e-01	40.9
DPM-dist2center	0.069	1.67e-04	41.3	SVM-FV	0.233	2.85e-38	49.1
Occlusion	0.163	3.17e-19	45.2	SVM-Att-Orac	0.150	1.90e-16	48.1
Aesthetic-Generic	0.139	2.29e-14	46.1	I2C-Att-Orac	0.126	5.52e-12	44.1
Aesthetic-Animal	0.185	1.73e-24	47.4	DAP-Orac	0.113	5.06e-10	44.2
Memorability	0.112	9.16e-10	43.7	DAP-Pred	0.063	5.95e-04	42.7

the SRC and AP measures, both of them lead to similar conclusions. We now provide a detailed discussion.

**Scene layout.** The strongest correlation is observed for the **BB-size** indicator. This shows that the object size plays a crucial role for image iconicity. We also inspected the distribution of the scores as a function of the BB size in our training set, and observed that the scores increase as the ratio between the BB area and full image increases up to around 0.85 and then decreases slightly. This may be because objects that are too big become less iconic as a certain amount of context is necessary. Also, the location predicted by the DPM detector allows producing a score that is almost as correlated with iconicity as the ground-truth BB, showing that a detector could successfully replace annotations in our scenario. The distance between the object center and the center of the image exhibits a lower correlation (see **BB-dist2center** and **DPM-dist2center**).

Considering the major role of object size, we could have applied indicators only to the object location (*i.e.* cropped images). We discarded this option as i) this is not consistent with the annotation acquisition (users were shown the full images), and ii) this applies a transformation to the images, which is different from evaluating the iconicity of an image as it is.

**Occlusion.** The correlation with the occlusion indicator is also high, confirming our intuition that the number of visible parts is related to iconicity. As expected, images of occluded objects are not iconic. We also trained a classifier on top of the binary part-visibility vectors to learn the importance of the parts, but could not improve the correlation.

**Aesthetic evaluation.** Both our aesthetic scores exhibit significant correlation with iconicity measures, proving the importance of this criterion. Even though it has been trained with much fewer ( $\sim$  two orders of magnitude fewer) images, the aesthetic model trained specifically with animal images performs better than the generic one (+0.046 correlation and +1.31 AP).

**Memorability.** The correlation with the memorability score is lower than with the aesthetic one, but a correlation is still observed. This implies that memorability and iconicity share common properties.

**Clustering.** Indicators computing a similarity between images and cluster centers of their class show a significant but small correlation with iconicity, which is surprising considering that it is the most commonly used indicator in the literature. The FV-based indicator (**Cluster-FV**) performs better than the GIST one (**Cluster-GIST**). This can be explained by the fact that the FV is a richer descriptor than the GIST.

**Object classifier scores.** We observe a higher correlation using the **SVM-FV** indicator. Based on the FV representation, a discriminative classifier trained to recognize a class can pretty well capture and predict iconic properties of that class, as already shown by [11] for different features. Note that this classifier is not perfect (our setting yields a top-1 classification accuracy of 32.8% on the corresponding 200-species classification problem) but is good enough to capture interesting properties. On the other hand, a SVM classifier trained with GIST descriptors produces a very low correlation. This is not surprising as the corresponding 200-class classification results drop to 4.7% top-1 accuracy.

**Attributes.** We observe that **SVM-Att-Orac** produces lower correlation than **SVM-FV**. We see two possible explanations. First, it could be better to let the classifier decide what a good representation of the class is from the image features (*e.g.* the FVs) themselves than using manually chosen attributes as a proxy. Second, attribute annotations could be noisy as each image was annotated by a single AMT worker [34]. The indicators **I2C-Att-Orac** and **DAP-Orac** perform well, but the best attribute indicator is **SVM-Att-Orac**.


**Correlation between indicators.** We also measure the correlation between the different indicators. SRC scores for a subset of the most promising indicators are presented in Table 2 (as color values). First, we observe the high correlations obtained by all methods using the class information. The two most correlated methods are the **SVM-FV** and the **Cluster-FV** (this pair obtains the very high value of 0.706), and the second highest correlation is the attribute method together with the **SVM-FV**. The aesthetic and the memorability indicators have very low correlation with the other methods. In the next section, we will show that these class-independent indicators nicely complement the class-based ones.

### 5.3 Prediction of image iconicity

From the previous study, we have identified indicators that best capture information related to image iconicity. We now show that they are complementary.

**Combination of the most relevant indicators.** In a first scenario, we assume that all annotations are available, and we combine the oracle versions of our indicators. We selected **BB-size**, **Aesthetic-Animal**, **Occlusion**, **Memorability**, **Cluster-FV**, **SVM-FV**, and **SVM-Att-Orac** due to their good correlation with iconicity. We also consider a second scenario, where annotations are replaced by predictions, except for the class label that is always considered as known (consistently with our scenario). We combine **DPM-size**, **Aesthetic-Animal**, **Cluster-FV**, **Memorability**, **SVM-FV**, and **DAP-Pred**. There is no occlusion criteria in this second scenario, as it seems unrealistic to train part detectors of a high enough quality to estimate the level of occlusion of the object. To make

**Table 2.** Spearman’s Rank Correlation (SRC) between each pair of selected indicators.



**Table 3.** Combinations of indicators to predict image iconicity. The SRC, corresponding p-value, and average precision (AP) are reported.

Method	SRC	p-value	AP
Combining Oracle Indicators			
Average	0.370	5.02e-98	58.0
SVM-bin	0.401	3.52e-116	60.9
SVM-rank	0.415	4.94e-125	60.4
Combining Predicted Indicators			
Average	0.303	1.55e-64	53.2
SVM-bin	0.353	2.23e-88	59.3
SVM-rank	0.350	2.39e-86	56.4
Direct Iconicity Predictors			
DIP-bin	0.372	7.34e-88	60.2
DIP-rank	0.375	1.08e-100	60.3

them comparable, the scores of each indicator are whitened (average 0, and std 1). We considered two combination methods. In the first one, we average all scores. In the second one, indicator scores obtained for our training images are used to train an SVM classifier that weights the different indicators. Both the binary and ranking SVMs are considered. All results are reported in Table 3.

First we observe that the average oracle indicator performs already significantly better than any indicator taken independently. This result shows that a class-independent combination works which means that we could predict iconicity for a new bird class. When learning a classifier, up to 2.9% of AP (0.045 SRC) can be gained. Second, we see that the predicted indicators (average and learnt) yield results that are quite close to the oracle ones, demonstrating the applicability of the combination without the optimistic oracle assumption. Qualitative results are presented in Fig. 3.

**Direct Iconicity Prediction.** The direct approach yields a high accuracy, comparable to our oracle indicators, showing that iconicity can be directly learnt from data. Yet, this indicator does not allow understanding what makes an image iconic. When combining the best discriminative approach (**DIP-rank**) and the best learnt oracle (**SVM-rank**), we obtain a correlation of 0.459 and an AP of 64.7. If we combine **DIP-rank** with the best prediction-based combination, we still obtain a 0.420 correlation and 63.7 AP which again shows that we could predict such iconicity score even on a scenario where fewer annotations are available. The best combination we obtain, 0.459 is quite close to 0.485 and 0.497, that correspond to the correlation between groups of annotators for the same images (discussed in section 3). This inter-person correlation gives an upper-bound on the maximum prediction accuracy we can achieve, and we observe that our best combination almost obtains this.

**Beyond birds.** We show qualitatively that our findings generalize beyond birds. For this purpose, we used images of the Fine-Grained Competition (FGcomp’13) [1] and applied directly the combination of iconicity predictors learned on the

bird images to two completely different objects: planes and shoes. Note that on these datasets, we did not have access to parts and attributes and therefore we could not use the **Occlusion** and **SVM-Att-Orac** predictors. The results are provided in Fig. 4 and show that the chosen most/least iconic images are plausible, although the system was trained on very different objects (birds).

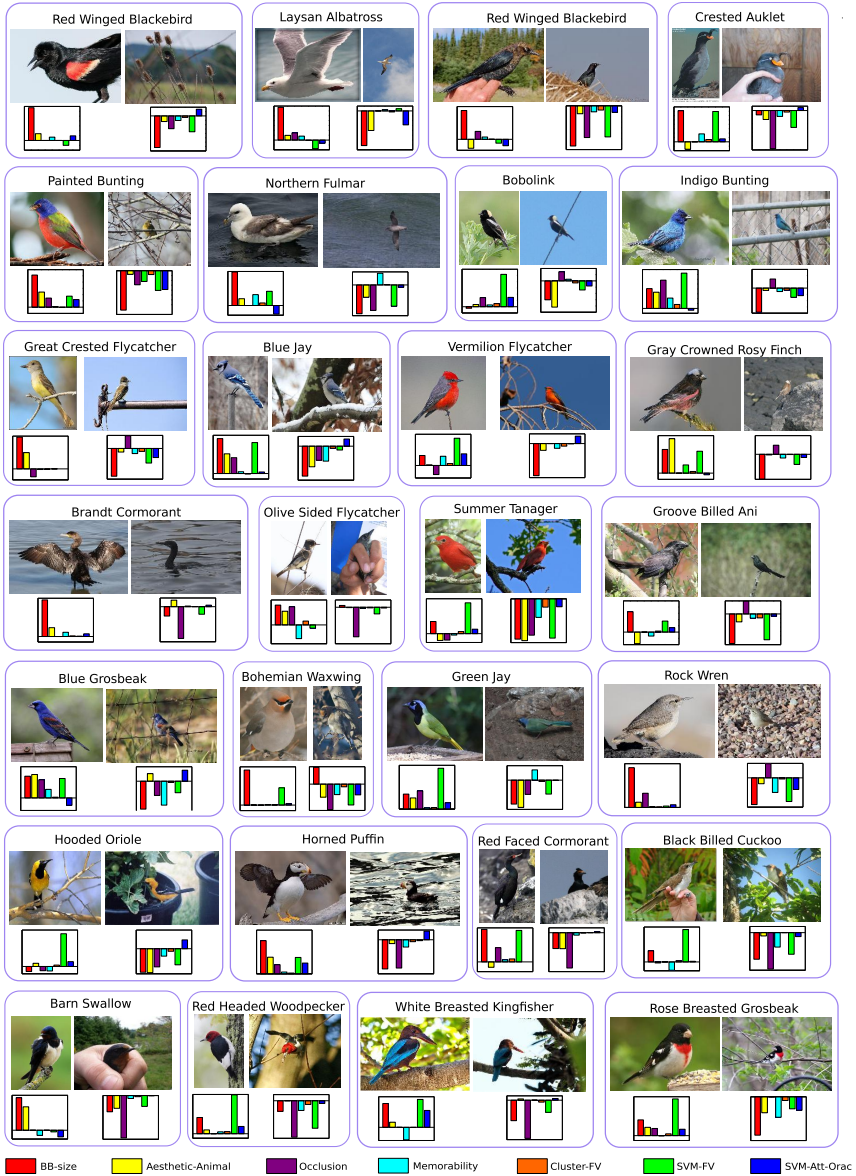
## 6 Conclusion

The goal of this study was to verify that iconicity is an image property rated consistently across multiple annotators and to understand what makes an image iconic. Toward this goal, we conducted an extensive study which involved collecting iconicity annotations from a set of users, and proposing a large set of possible properties that can predict iconicity. These properties cover all the ones considered in previous work, plus new ones we proposed. Some are class-dependent, and some are totally generic. Our study showed that in the fine-grained context, these properties used with the right descriptor are all relevant to iconicity, and are complementary with each other. We also proposed to directly predict iconicity with a classifier discriminatively trained on iconicity rating, and showed the good performance of this simple yet novel scenario. We expect all these findings to be useful in many computer vision applications of practical value.

## References

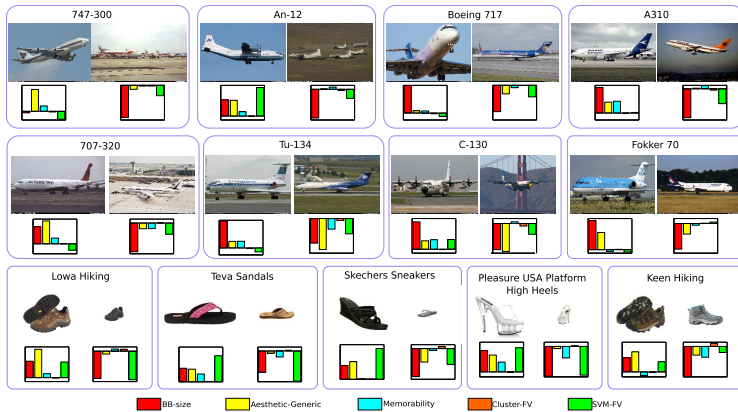
1. FGComp'13. <http://sites.google.com/site/fgcomp2013/>
2. Berg, T., Berg, A.: Finding iconic images. In: IVW at CVPR (2009)
3. Berg, T., Forsyth, D.A.: Automatic ranking of iconic images. Tech. rep., U.C. Berkeley (2007)
4. Blanz, V., Tarr, M., Bülthoff, H., Vetter, T.: What object attributes determine canonical views? Tech. rep., MPI (1996)
5. Branson, S., Wah, C., Babenko, B., Schroff, F., Welinder, P., Perona, P., Belongie, S.: Visual Recognition with Humans in the Loop. In: ECCV (2010)
6. Bülthoff, H., Edelman, S.: Psychophysical support for a two-dimensional view interpolation theory of object recognition. PNAS (1992)
7. Caruso, J., Norman, C.: Empirical size, coverage, and power of confidence intervals for spearman's rho. Educational and Psychological Measurement (1997)
8. Clinchant, S., Csurka, G., Perronnin, F., Renders, J.M.: XRCE's participation to ImageEval. In: ImageEval Workshop at CVIR (2007)
9. Denton, T., Demirci, F., Abrahamson, J., Shojoufandeh, A.: Detecting canonical views for view-based 3-D object recognition. In: ICPR (2004)
10. Ehinger, K., Oliva, A.: Canonical views of scenes depend on the shape of the space. Cognitive Science Society (2011)
11. Ehinger, K.A., Xiao, J., Torralba, A., Oliva, A.: Estimating scene typicality from human ratings and image features. Cognitive Science Society (2011)
12. Farhadi, A., Endres, I., Hoiem, D., Forsyth, D.: Describing objects by their attributes. In: CVPR (2009)
13. Felzenszwalb, P.F., Girshick, R.B., McAllester, D., Ramanan, D.: Object detection with discriminatively trained part based models. TPAMI (2010)

14. Ferrari, V., Zisserman, A.: Learning visual attributes. In: NIPS (2007)
15. Garg, S., Berg, T., Mueller, K.: Iconizer: A framework to identify and create effective representations for visual information encoding. In: Smart Graphics (2011)
16. Girshick, R., Felzenszwalb, P., McAllester, D.: Discriminatively trained deformable part models, release 5. <http://people.cs.uchicago.edu/~rbg/latent-release5/>
17. Isola, P., Xiao, J., Torralba, A., Oliva, A.: What makes an image memorable? In: CVPR (2011)
18. Isola, P., Xiao, J., Parikh, D., Torralba, A., Oliva, A.: What makes a photograph memorable? TPAMI (2013)
19. Jing, Y., Baluja, S., Rowley, H.: Canonical image selection from the web. In: CIVR (2007)
20. Joachims, T.: Optimizing search engines using clickthrough data. In: SIGKDD (2002)
21. Ke, Y., Tang, X., Jing, F.: The design of high-level features for photo quality assessment. In: CVPR (2006)
22. Lampert, C.H., Nickisch, H., Harmeling, S.: Learning to detect unseen object classes by betweenclass attribute transfer. In: CVPR (2009)
23. Lazebnik, S., Schmid, C., Ponce, J.: Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In: CVPR (2006)
24. Li, X., Wu, C., Zach, C., Lazebnik, S., Frahm, J.M.: Modeling and recognition of landmark image collections using iconic scene graphs. In: ECCV (2008)
25. Lowe, D.: Distinctive image features from scale-invariant keypoints. IJCV (2004)
26. Mezuman, E., Weiss, Y.: Learning about canonical views from internet image collections. In: NIPS (2012)
27. Murray, N., Marchesotti, L., Perronnin, F.: Ava: A large-scale database for aesthetic visual analysis. In: CVPR (2012)
28. Oliva, A., Torralba, A.: Modeling the shape of the scene: A holistic representation of the spatial envelope. IJCV (2001)
29. Palmer, S., Rosch, E., Chase, P.: Canonical perspective and the perception of objects. Attention and performance (1981)
30. Raguram, R., Lazebnik, S.: Computing iconic summaries of general visual concepts. In: Internet Vision Workshop at CVPR (2009)
31. Raguram, R., Wu, C., Frahm, J.M., Lazebnik, S.: Modeling and recognition of landmark image collections using iconic scene graphs. IJCV (2011)
32. Sánchez, J., Perronnin, F., Mensink, T., Verbeek, J.: Image classification with the fisher vector: Theory and practice. IJCV (2013)
33. Simon, I., Snavely, N., Seitz, S.: Scene summarization for online image collections. In: ICCV (2007)
34. Wah, C., Branson, S., Welinder, P., Perona, P., Belongie, S.: The CUB-200-2011 Dataset. Tech. rep., CalTech (2011)
35. Wah, C., S.Branson, Perona, P., Belongie, S.: Multiclass recognition and part localization with humans in the loop. In: ICCV (2011)
36. Weinshall, D., Werman, M., Gdalyahu, Y.: Canonical views, or the stability and likelihood of images of 3d objects. In: Image Understanding Workshop (1994)
37. Weyand, T., Leibe, B.: Discovering favorite views of popular places with iconoid shift. In: ICCV (2011)



**Fig. 3.** For a set of representative bird species, images with the highest and lowest predicted iconicity based on our learnt combination of oracle indicators are shown respectively as the left and right images of each box. Below images, histograms display the contribution of each indicator in the decision process. For instance, for the “Brandt Cormorant”, the **BB-size** plays a crucial role in deciding the choice of the most iconic image while for the “Summer Tanager” almost all indicators play a role in the choice of the least iconic image (best viewed in color).





**Fig. 4.** For a set of representative classes extracted from two subsets of the FGcomp'13, namely shoes and planes, and for the available indicators, this figure shows images with the highest and lowest predicted iconicity based on our learnt combination of oracle indicators, respectively as the left and right images of each box. Below images, histograms display the contribution of each indicator in the decision process (best viewed in color).